

S 1

NO. 007

Why OpenRouter Could Be the Next Great Platform

I think of OpenRouter as a great business.

JULY 17, 2026

I think of OpenRouter as a great business. It aggregates practically every Large Language Model, then bundles access and takes a flat fee. Their main value proposition for an enterprise is cost savings. They give you clear access to every model and allow you to judge the optimal use for your business. It makes accessing models extremely easy and developer-centric, yet jeopardizes the position of the frontier labs.

Today, The Information broke news of acquisition interest in OpenRouter. Shocker: while I can speculate on who might be interested, what is clearly more interesting is why anyone would want to buy OpenRouter. here

Simply for two reasons: firstly, it has clear platform dynamics. It is an aggregator of AI models, and once enterprises use it, that creates switching costs since all your AI expense is in one place: logs, budgets, etc, so switching would be painful and generally not make a lot of sense.

That's clear, hence the 1.3 billion valuation. But it's interesting for the interested acquiring parties because OpenRouter threatens to commoditize Anthropic and OpenAI. For most enterprises, the differentiation between these models is thin, and considering that yesterday Kimi K3 beat a lot of benchmarks and for a fraction of the cost, albeit speculated distillation, it further proves why OpenRouter should be a dominant business.

So perhaps a frontier lab is hoping to acquire them to defend themselves from becoming a utility. Or maybe Google, which led OpenRouter's 113m series B. So it would likely be a big mistake to sell, considering the amount of leverage the company currently has.

The reason a commodity business would be bad for a frontier lab is that the fixed cost investment of training needs to be spread across as many enterprises as possible, yet if a platform has hundreds of models and many, if not most, being cheaper, it's hard to command any pricing power at all. And all the capital invested in these labs would essentially be useless since growth would falter, and the current valuations, which are way out of touch with reality, would crash.

In terms of valuation, their latest round was in May at 1.3 billion. Given the market and the dominant position of the company, 40 to 50 x their revenue, which was last stated at somewhere around 50 million. So in the range of 2 billion. Yet if they can aggregate the demand of the enterprise world and take a comfortable 5.5% on every company's model usage, OpenRouter could become a cash-printing machine, far more value-accretive than a measly 2 billion exit.

So the recent Stripe acquisition talks around the 10 billion number seem expensive, based on their recent 140 million in ARR. But they are growing rapidly; their revenue was 40 million in November, so around a 17% monthly growth rate. Here is a conservative revenue forecast -

OpenRouter — Projected Financials					
Year	Revenue (ARR)	Growth	Gross Profit (85%)	Operating Income (40%)	Free Cash Flow (35%)
2026	\$140M	—	\$119M	\$56M	\$49M
2027	\$280M	100%	\$238M	\$112M	\$98M
2028	\$490M	75%	\$417M	\$196M	\$172M
2029	\$784M	60%	\$666M	\$314M	\$274M
2030	\$1.14B	45%	\$969M	\$456M	\$399M
2031	\$1.54B	35%	\$1.31B	\$616M	\$539M

This assumes that OpenRouter becomes a great enterprise software company, but not the dominant platform it appears to be becoming. Revenue is likely 3 to 4 billion in that scenario. Plus, Stripe is not trying to maximize FCF; rather create strategic enterprise lock-in and layers on additional services. It also views OpenRouter as a dominant player in the AI era, and can use it as a beachhead to add on Stripe services.

